Initial Draft: August 2022

PRELIMINARY: COMMENTS WELCOME

AGING IN STYLE:

Does How We Write Matter?

Daniel S. Hamermesh and Lea-Rachel Kosnik*

ABSTRACT

The scholarly impact of academic research matters for academic promotions, influence, relevance to public policy, and others. Focusing on writing style in top-level professional journals, we examine how it changes with age, how stylistic differences and age affect impact, and how style and prior scholarly output relate to an author's subsequent achievements and labor-force decisions. As top-level scholars age, their writing style begins to differ from others'. The impact (measured by citations) of each contribution decreases, due to the direct effect of age and much smaller indirect effects through style. Scholars produce less top-flight work as they age, especially those who have produced less recently, whose work is less cited, and whose styles have been more positive. Previously less productive authors are more likely to retire.

*Hamermesh: University of Texas at Austin, IZA, and NBER; Kosnik: University of Missouri—St. Louis. We thank Shira Itzkowitz for extremely careful research assistance and Pavel Jelnov for providing some of the data on authors' ages. In the interest of transparency, we note that one of the authors is included in the data set.

They shall bring forth fruit in old age [Psalm 92:14]

I. Introduction

The essence of academic scholarship is contained in what academics write, and the rewards to successful writing - research that affects the public and other scholars - are substantial. These range from purely monetary (as an immense literature—with early examples of Holtmann and Bayer, 1970, in the natural sciences; Hamermesh *et al*., 1982, in economics; Diamond, 1986, in mathematics, shows), to honors ranging from appointment as a Fellow in some academic society to the pinnacles—a Nobel Prize in the natural sciences (and economics), a Fields Medal in mathematics (Hamermesh and Pfann, 2012; Borjas and Doran, 2015), and others. Our first major question goes behind the effects of successful scholarship to ask: Does **how** we write affect the success of our writing? Before we can answer that question, however, we examine how our writing styles vary with our demographic characteristics, most importantly, our age/experience as researchers. After all, if academic success is related to both age and the style in which we present our research, we need to separate out the indirect effects of age through style to the direct effects of age on success. Parsing out these causes allows us to get a glimpse into one possible source of the well-recognized decline in creative activity with age (Lehman, 1953; Levin and Stephan, 1991; Weinberg and Galenson, 2019; and many others), which is our second major focus.

We can delve further into the causes of the relatively short scholarly lives of academics by obtaining information on the ages when they stop producing fundamental work. Do they retire from academia, so that the incentive to produce such work is diminished, perhaps to zero? If so, do they leave because their academic production has already substantially slowed? Has their production not slowed, but become less well recognized by their peers? Or have they simply tired out or become technologically obsolete, which also possible residual explanations? We cannot test or rule out all explanations, but with information on the life-cycle patterns of academic production and its success, we can consider some of them.

To answer these questions, we need information on publishing patterns over scholars' lives, on the style of their publications, and on the impact of their research on the scholarly community, all in relation to the author's age at which the research appears. In order to examine possible reasons for declining activity, we also need information on the age at which those in our sample stopped producing top-level research. These data requirements begin to be satisfied in a sample of all publications in the so-called "Top 5" economics journals that appeared between 1969 and 2018. This sample contains the most influential economics publications over the past half century, and the individuals in the sample are the upper crust of contributors to economic knowledge. Using textual analysis to measure style, for each article we then obtain its subsequent citations to measure its scholarly impact. For each author we obtain her/his date of entry into academe, and for each we acquire information on whether they remained unretired in academe in 2018.

Section II details the sample we construct, describes the measures of sentiment, and provides descriptive statistics of the publications. In Section III we present the first set of main results, linking style to age and describing how deviations in style from norms that prevailed at the time of publication and in the sub-field of the research vary with age. Section IV examines how style and age relate to citations and teases out the direct and indirect (through style) impacts of age on the subsequent impact of scholarship. Finally, in Section V we analyze whether and how publishing success and the style of published research affect the likelihood of continuing to produce top-notch academic research, thus getting at some of the possible explanations of age-related slowdowns in productivity.

II. The Sample, and the Measurement of Sentiment

A. Publications in Economics, 1969-2018

The corpus of texts that we analyze consists of all 16,827 research articles published in English in the "Top 5" economics journals: *American Economic Review (AER), Econometrica (ETRCA), Journal of Political Economy (JPE), Quarterly Journal of Economics (QJE), and Review of Economic Studies (REStud),* from 1969-2018.[1] Entries not included in the dataset are editor's notes, conference announcements and programs, auditor's reports, indexes, other similar non-research focused entries, and articles in the AEA *Papers and Proceedings*. Special symposium articles are included. Importantly, the

[1]Some of these journals, especially in earlier years, included an occasional article in French or German.

dataset utilizes entire articles, and not just article titles or abstracts, as is sometimes the basis of corpora in the literature that investigates academic research.

We exclude all entries that are comments/replies/rejoinders, etc., and also those that are Nobel or presidential addresses (American Economic Association or Econometric Society), since the former may depend on the original article being discussed, while the latter need not be purely scientific articles. These exclusions reduce the sample to 15,138 articles. With multiple authors on a majority of these articles, we have over 20,000 author/article entries. Many of the authors are “one-hit wonders,” and many others appear only a few times. Since we wish to concentrate on the life-cycle relationship of age to style and scholarly impact, we restrict the sample to authors with at least five articles among the 15,138. For each of these highly successful authors we attempted to obtain the year when they began their careers, which we take as the year of receipt of the Ph.D.[2] Through online searches and emails, both to authors and, where necessary, their colleagues, we obtained this measure for all but one individual (who authored six articles in the sample). Our final sample thus contains 12,814 articles authored by 1,389 different individuals.[3] We also record the gender of each author (since Kosnik, 2022, demonstrates gender differences in style even within the same sub-field in economics).[4]

Of particular interest is the cohort of individuals who entered the profession (received their doctorate or equivalent) between 1969 and 1978. For these 359 scholars (which we call the 1970s cohort), who authored 3,562 of the articles in the sample, we can observe nearly their entire professional careers, thus creating a longitudinal sample of the leading scholars in this cohort whose members have had at least 40 years to publish scholarly research.

[2]For the 0.5 percent of authors without a Ph.D. we add five years to the date when they received an undergraduate degree and count their professional experience from that year.

[3]Because several articles may have the same pair, triplet, or even quadruplet of co-authors in the sample, only 9,280 separate articles are included. In calculating sample statistics describing authors and in estimating models, we thus weight each of the 12,814 observations by the inverse of the number of times it appears in the sample.

[4]In terms of data collected, the most similar study is Coupé et al., (2006), which considered totally different issues from the aging and style questions examined here.

The main source of sample selectivity is along the criterion of scholarly success—having at least five research articles in these most visible scholarly outlets in economics. We recognize that the "Top 5" are only a few of the 182 economics journals that were indexed in *EconLit* in 1969 (and of the more than 1,000 included there today), and that many articles in other outlets receive more attention (Oswald, 2007; Heckman and Moktan, 2020). On average, however, articles published in these journals do attract the most attention (Hamermesh, 2018). The exclusion of authors with few "Top 5" publications is restrictive, but it allows us to follow careers over a reasonably long period of time. We admittedly concentrate on the careers of academic stars, so that in none of our analyses can we infer anything about the careers of scholars with relatively top-level scholarly contributions.[5]

For each entry we have its length in pages (which we normalize to the word count of the *AER* before 2000). Since styles may differ by type of article, we also include the first-listed top-level *JEL* classification of each article (*JEL* Codes A-R and Z). These are aggregated into five groups: Theory and methodology (*JEL* = C); microeconomics and industrial organization (*JEL* = D, L); macroeconomics, international economics, and financial economics (*JEL* = E, F, G); public economics, health/education, and labor and demographic economics (*JEL* = H, I, J); and other. We also know the decade of publication, 1969-78, …, 2009-18, which we use in transforming the raw measures of sentiment that we create.

The top panel of Table 1 describes the characteristics of the articles. They are distributed fairly evenly across the five *JEL* groups, with the exception of the smaller category of other—miscellaneous—articles. Despite the apparent growth in publishing, the distribution of articles is nearly uniform across the decades. The growth in publishing is explained by the logorrhea of authors publishing in these journals, a near tripling of page lengths over the five decades.

As is well-known, publishing top-level economics is a young person's game (see Hamermesh, 2013, for cross-section evidence), as the kernel density estimates for the entire sample in Figure 1a

[5]During the decade of the 1970s perhaps 8,000 Ph.D. degrees in economics were conferred in the U.S. Our sample of 359 usable observations thus represents the most successful five percent of publishers in that cohort.

demonstrate and as shown by Figure 1b for the longitudinal data describing the 1970s cohort.[6] In this cohort the median age post-Ph.D. at an article's publication is 10 years, with only 1.2 percent of articles published before an author received his/her doctorate. Among those with Ph.D. degrees received before 1972, only 6 percent of the articles they published in this half-century were written when they were more than 35 years post-Ph.D. Only 0.7 percent of the articles published in 1969-78 contained female authors who were in the sample, a percentage that reached 5.7 in the decade 2009-18; and the 1970s cohort is only 1.2 percent female, while 7.0 percent of sample members with doctorates 1979 or later are women.

The second panel of Table 1 describes the achievements of this selected sample. Nearly one-quarter of the 1,389 authors barely qualified for inclusion, with only five papers published in these outlets. The maximum number of articles anyone published in these journals during this half century is 60. Restricting the calculations to the 1970s cohort, the distribution looks quite similar to the overall distribution, although it is shifted slightly to the right.

B. *Measuring the Sentiment of Economic Research*

Sentiment analysis is a technique for identifying the emotive tenor of a piece of writing. The use of sentiment measures in economics is discussed by Gentzkow et al. (2019), and they have been used in many areas of scholarly research, including as examples analyses of the Old and New Testaments (Houk, 2002; Kenny, 1986), and examinations of the authorship of the individual Federalist papers (Mosteller and Wallace, 1963). We utilize three sentiment scores in this research: A positive/negative sentiment score (POSN), a certainty/tentativeness sentiment score (CERT), and a contemporary/past sentiment score (CONP). Each score j has been determined as a net count of all relevant word or word-phrases in document i divided by the total number of relevant words:

(1) $$z_{iaj} = \Sigma c_{iaj} - \Sigma t_{iaj} / \Sigma(c_{iaj} + t_{iaj}) \ ,$$

[6]This is the largest cohort in the sample, accounting for 26 percent of authors. The pre-1969 cohort included 14 percent of authors, the 1979-88 cohort 23 percent, the 1989-98 cohort 20 percent, and two youngest cohorts together 17 percent. The year of receipt of Ph.D. ranged from 1937 to 2014.

j=1..,3, and where z_{ij} is the net score for article i by author a along criterion j. c_{iaj} is the count of its positive (certain) words, and t_{iaj} the count of its negative (tentative) words. CONP is calculated based on the c_{ia3} indicating future or present tenses in verbs, t_{ia3} indicating verbs in the past tense. Each of the three indicators is thus based on counts of words classified into two contrasting types.

If POSN >(<) 0, we infer that an article has a net positive (negative) emotive tone. The size of the final score indicates the degree of its net positivity or negativity. If CERT >(<) 0, an article has an overall emotive tone of certainty (tentativeness). If *CONP* >(<) 0, the article has a contemporary (past-focused) emotive tone. For all three measures the size of the sentiment score indicates the degree of the overall sentiment.

Key to any sentiment score are the words and phrases that comprise the c_{ij} and t_{ij}. Appendix Tables A1-A3 provide examples of the kinds of words and phrases in each of the three sentiment scores. The dictionaries utilized for this analysis were built up from the Harvard IV dictionary (https://textanalysis.info/pages/category-systems/general-category-systems/harvard-iv-dictionary.php), the Linguistic Inquiry and Word Count (LIWC) dictionary (Pennebaker, 2015), and the Regressive Imagery dictionary (Martindale, 1990), with edits made to fit the context of writing in academic economics.[7] These edits primarily involved recognizing econometric-based words as neutral, rather than as indicative of emotive content. For example, "average," "limit," "regression," "subtract," and "ordinary" were marked as indicative of negative sentiment in the original dictionaries, but were made neutral for analyzing economics articles. Similarly, "aggregate," "natural," "validity," "append," and "value" were all marked as indicative of positive sentiment in the original dictionary, but were also treated as neutral for this analysis. Dictionary creation is a somewhat subjective endeavor, which is why we relied, as much as possible, on the category

[7]Tailoring the dictionary to the context is important, as some words have different meanings in different contexts. "Vice," for example, would be categorized as a negative word in most situations, but in a human resources managerial handbook it might refer primarily to vice-presidents and so be categorized in that context as neutral. It then would have no bearing on such a handbook's positive/negative sentiment score.

dictionaries created by previous researchers which have been honed over many years of use. We tailored them only (as is standard in the literature) within the specific context of economics and econometrics.[8]

Each of the articles in the corpus was entered into a relational database where variables associated with the articles could be independently analyzed, for examples, year of publication, journal of publication, page length, and of course, author's age. The text itself was left unstructured and organized within a vector-space model (VSM), where each element of the vector indicates the occurrence of a particular word or phrase within the paper. The vector elements were not transformed or weighted in any way, instead being left as raw frequency counts, so that if a given word was used more than once in a paper, its degree of emphasis was reflected in a higher count and thus a higher sentiment score.

The textual analysis yielded the measures z_{iaj}. Because there are trends in style (Kosnik, 2022) and differences in style across sub-fields, we transform each z_{iaj} as:

(2) $z^*_{iaj} = z_{iaj} - z'_{..j}$,

where $z'_{..j}$ is the score averaged over all articles by all authors in a *JEL* group in a decade (so that each score is adjusted by the norm of sentiment for its subfield and time). The calculations of the z^*_{iaj} also allow examining the size of the departure of style, whether positively or negatively, from the subfield/time norm describing the article, a departure which we measure as z^{*2}_{iaj} . Like the measures of sentiment themselves, these departures may be related to the authors' ages and to the success of their articles.[9]

The bottom panel of Table 1 lists summary statistics of the z_{ij}, z^*_{ij}, and z^{*2}_{iaj}. On average the sentiment of the articles in our sample is quite negative, they are written in a very tentative voice, but they do tend to be contemporary oriented. Sixteen percent of the articles have a net positive sentiment, and four

[8]After the initial word counts and sentiment scores were calculated, spot checking with KWIC (keyword-in-context) was performed to make sure the words being categorized as negative or positive really indicated such sentiment in the article.

[9]Using the quadratics in z^*_{iaj} to measure departures from norms is arbitrary, implicitly assuming increasing effects as the departure increases. We re-estimated all the models in Sections III and IV, replacing z^{*2}_{iaj} by $|z^*_{iaj}|$. The coefficient estimates become slightly less significant, but, more important, the fits are not as good. This suggests that the implicit assumption of increasing effects regardless of the sign of the departure from the norm underlies the data.

percent express certainty in their sentiments, although almost none contain a net past-oriented sentiment. The crucial point to note for our empirical analyses is that there is substantial variation in sentiment across the sample along all three criteria.[10] Moreover, as Appendix Table B2 shows, while the correlations among the three measures of sentiment, the deviations and their squares, across the samples are positive, they are not very large. The three measures of sentiment in an economics article are nearly independent.

III. Age, Style, and Style Norms

We first consider how style and style norms relate to age nonparametrically by examining the local polynomial smoothed relationship between a sentiment measure and the Ph.D. age of authors at the time their paper was published. Figures 2a-2c show these for each of POSN, CERT, and CONP, including 95-percent confidence bands around the estimates. While these figures cover the entire range of ages when the author's articles appeared, the paucity of publications before an author received his/her Ph.D., or after Ph.D. age 35 unsurprisingly makes the confidence bands over those ranges very wide. The most useful comparisons are of the patterns of sentiments when the authors are between Ph.D. ages 0 and 35; assuming a Ph.D. is received at age 28, roughly between actual ages 28 through 63.

These comparisons demonstrate a monotonic and highly statistically significant increase in the positivity of writing style with age over the relevant Ph.D. age range (Figure 2a). Conversely, there is a significant monotonic decrease in the certainty of writing styles over this age range (Figure 2b). There is essentially no relation between age and the present/past orientation of the authors' styles (Figure 2c), except that even with the small sample of very senior authors, there is a significant decrease in present orientation after a Ph.D. age of 35.

While allowing a function-free view of the sentiment-age relationships, the estimates in Figures 2a-2c cannot allow for other characteristics (of authors and articles) that might determine the style in which the articles are written. The top panel of Table 2 presents simple linear estimates relating the deviation in

[10]There are also significant differences across the five journals, with all of them being more positive and more contemporary-oriented than the *AER*, and all but the *QJE* being written in a more certain voice than the *AER*, as Appendix Table B.1 shows.

sentiment (the z^*_{ij}) to Ph.D. age, but holding article length constant. (The estimates shown in the Table are the raw estimates multiplied by 100.) The standard errors of the parameter estimates are clustered on authors. The estimates essentially reproduce the results in Figures 2a and 2b. For CONP, however, the linear estimate yields a significant decline in present orientation with age of the entire sample, no doubt because of the significant sharp drop observed in Figure 2c among the oldest authors.

With the demonstration in Appendix Table B1 that there are differences in style across journals, and knowing that using deviations in sentiment may not fully account for trends and sub-field differences, the second panel in Table 2 adds as controls journal indicators, *JEL* group, decade of publication, and author's gender. These controls produce almost no changes in the estimated relationships of the style measures to Ph.D. age, with sentiment becoming more positive, less certain, and more past-oriented with age. The final panel includes author fixed effects, thus adjusting for any personal idiosyncrasies in style. The signs of all three estimates remain the same, with the impacts of positivity (present orientation) remaining statistically significantly positive (negative).

The right-hand side of each panel in Table 2 presents the same estimates but only including authors in the 1970s cohort. This restriction allows concentration on a group whose backgrounds and professional life experiences were probably more homogeneous than those of the entire sample. The estimates in the first two panels are similar in magnitude in most cases to those for the entire sample, although with a sample size only 28 percent of that in the entire group, their standard errors are larger. The fixed-effects estimates are much smaller than those for the entire sample, but they still show the positive positivity-age relationship, and the negative relationships of the other two sentiments to age. The overall conclusion from Figures 2 and Table 2 is that there is some evidence that sentiment changes, all else equal, as authors continue writing, becoming more positive and less present-oriented (opposite from the secular trends for the profession as a whole found in Kosnik (2022)).

With co-authorship increasing steadily over the half-century of our sample, perhaps the results simply reflect correlations of the number of co-authors with style. Adding the number of co-authors to the models in the lower two panels on the left-hand side of Table 2 hardly changes the estimated effects of author's age on writing style. Additional authors do, however, make the style more certain and less present-oriented. Adding the same measure to the estimates based on the 1970s cohort has similar effects.

One might be concerned that, with so many authors having only five entries in the sample, the results arise from the characteristics of the least successful among this group of very successful scholars. As robustness checks, we re-estimate the equations discussed above, first restricting the sample to exclude the 24 percent of authors (13 percent of articles) with "only" five publications, then excluding the 74 percent with fewer than 10 publications (46 percent of articles). The results with the first exclusion yield uniformly larger (in absolute value) effects than those shown for the entire sample in the second panel of Table 2. With the even stricter exclusion the effects of age on the deviations of style from norms become even larger, perhaps because the most prolific authors, those with ten or more publications in the sample, stake out their stylistic identities earlier in their careers than other authors.

Figures 3a-c show local polynomial smoothed representations of the relation between age and the z^{*2}_{iaj} –the squared deviations of the sentiment measures from their decadal/sub-field norms. The results are even clearer than in Figures 2: Deviations from the norms of positivity fall with age; those with certainty rise with age, while there is no relation of the squared deviations of present-orientation to age over most of the range (although the squared departures fall significantly within the small sample of very senior authors).

Table 3 presents the same models as in Table 2, with the same additional controls and the same sample restrictions, describing the determinants of the z^{*2}_{iaj}. The estimates for the entire sample vary depending upon the control variables included; but the fixed-effects estimates, which include the decade, journal, and *JEL* controls, demonstrate that increasing age leads to significant increases in the departure of style from decadal/sub-field averages. Once we account for author fixed effects, we observe that, as authors age, their writing increasingly differs from others working at the same time and in the same areas—they become more unusual. Restricting the sample to the 1970s cohort strengthens this conclusion: The estimates for all three measures of sentiment depart increasingly and significantly from the norm as authors age.[11]

[11]As with the estimates in Table 2, we examine the robustness of the estimated effects on the z^{*2}_{iaj} by adding a measure of the number of coauthors on each article. These additions do not alter any conclusions about the relationship between

The writing of scholars in economics increasingly differs—both more or less positively, both more or less certainly, and both more or less present-oriented—as they gain experience.[12]

Most of the estimated impacts of age on the deviation of sentiment and the squared deviation from the norm are statistically significant. In relative terms they range from the very small—a decrease of 0.02 standard deviations in the squared deviation of POSN in response to a two standard-deviation change in author's age, to fairly large—an increase of 0.13 standard deviations in the squared deviation of CONP, and an increase of 0.25 standard deviations in the squared deviation of CERT in response to the same change in age. Age is related to sentiment—significantly so—and the impact of age on the size of the departure, positive or negative, from decadal/sub-field norms is not small.

IV. Age, Style, and Citations

We measure the scholarly impact of articles by the number of subsequent citations received. We recognize the imperfections in this measure, but: 1) It is a relatively objective measure; 2) It correlates well with various outcomes, including salaries and departmental/institution rankings (Hamermesh, 2018); 3) It correlates well with subjective evaluations by teams of economists (Checchi *et al*., 2021); and 4) Although imperfect, citations are the standard metric used in the literature describing academic contributions. Ideally, we would use citation counts from the Web of Science, but we were only able to obtain them for 63.9 percent of the 12,814 observations. Accordingly, we obtained citation counts from Google Scholar for another 35.5 percent.[13] All the citation data are cumulative through August 2021.

---

age and style. Having additional coauthors is, however, associated with greater departures from all three style norms. We also examine the robustness of the estimated effects to restrictions on the sample by excluding the less successful members of the sample. We impose successive restrictions on the sample, initially excluding authors with only 5 entries, then those with fewer than 10 entries. Examining only the fixed effects estimates for the 1970s cohort, the former restriction does not qualitatively alter the results, while the more severe restriction increases the absolute values the point estimates for POSN and CONT, reduces that for CERT.

[12]We re-estimated all the models discussed in this section replacing indicators of the *JEL* group with the raw *JEL* classifications, and replacing the decadal indicators with the indicators of the year of publication. Neither of these changes altered the general conclusions drawn from the estimates shown in the tables.

[13]We were unable to obtain citation counts for 0.6 percent of the observations and hence exclude them from the analyses in this section, resulting in a usable sub-sample of 12,740 observations (and the same 1,389 distinct authors).

Google Scholar is much less restrictive than the Web of Science (Hamermesh, 2018). The average citation count of the one-third of the sample with data from the former source is thus 677 (s.d. = 2,082), while that from the latter source is 199 (s.d. = 440). Authors of articles whose citations are from the Web of Science are much younger than other authors (average age since Ph.D. 10.1 versus 14.9), which results from the fact that younger authors have published more recently and that articles with Google Scholar citation counts are disproportionately (97 percent) from the earliest three decades of our sample.[14]

Consider first the local polynomial fits of citations to Ph.D. age, shown in Figure 4a for the entire sample and Figure 4b for the 1970s cohort. While the relationship is very imprecise at the extremes of Ph.D. age (below age 0 and above age 35) because of the paucity of data in those age ranges, across the bulk of observations there is a clear negative relationship between citations and Ph.D. age. Moreover, the decline with age is not insubstantial; Figure 4a shows that going from age 0 to 35 cuts the estimated citations to an article by nearly half.

The difficulty with these figures is that they cannot account for the differences in the source of the citation data by year of publication (and thus implicitly by author's Ph.D. age), nor for growth in the number of journals citing economics articles, nor for the length of time over which a study could accumulate citations by August 2021. We thus estimate models similar to those presented in Tables 2 and 3, including in all the equations each author's Ph.D. age and a measure of sentiment (and controlling in each equation for the number of equivalent pages and an indicator for whether the citation measure is from the Web of Science or Google Scholar). Because the distribution of citations is highly skewed, all estimates are produced using median (LAD) regressions and, as before, standard errors are clustered on authors.

The top panel of Table 4 presents the estimated impact of sentiment and age on subsequent citations, first for the entire sample (left) and then for the 1970s cohort (right). Greater age directly reduces subsequent citations, as implied by the Figures. Also, however, articles written in more positive, more certain, and more contemporary styles generate fewer citations. For the entire sample all the impacts are

[14]All the results in this section for the entire sample were reproduced on the separate samples with Web of Science data or Google Scholar data, with no departures from the results tabled here.)

highly significant statistically, while even for the much smaller cohort sample they are at least nearly statistically significant at standard levels.

The middle panel of the table adds controls for journal, *JEL* group, and for each year of publication, thus accounting for all the difficulties noted about the local polynomial fits. The results remain essentially the same as in the top panel: A direct negative effect of age, and negative effects of more positive, more certain, and more contemporary styles on citations. With the addition of these controls, all the estimates for the entire sample remain highly significant, as do the direct effects of age in the 1970s cohort. The estimated impacts of sentiment in the1970s cohort remain negative, but only that of CERT is statistically significant.

The bottom panel in Table 4 includes the same variables as the middle panel, but includes all three sentiment measures in the same equation. As such, we view these estimates as the most reliable. Given the low correlations among the sentiment measures, these pooled estimates reproduce the results in the middle panel. We can conclude that articles written by older authors are cited less, all else equal—a direct negative effect of age on scholarly success. Since we showed that age affects style, the negative (for CERT statistically significant) impacts of style on citations reflect the indirect effect of age on citations. We aggregate the three effects—summing up the indirect effects of age through all three sentiment measures—below.

While positive deviations of all three measures of sentiment reduce citations significantly or nearly so, the more important question is how large these reductions are. Taking two-standard deviation increases in sentiment scores, we calculate, based on the pooled equation shown at the bottom of Table 4, that these increases reduce citations by 9 percent, or 0.02 standard deviations of the citation measure. Writing in a more positive, more certain, or more present-oriented way reduces the scholarly impact of one's articles.

Table 5 produces analogous results, but for the z^{*2}_{iaj}, the squared deviation measures, which we present first with no controls (other than page-length and the source of citations), then in the middle panel with large vectors of controls, then in the final panel pooling the sentiment measures into one equation. Given the low intra-correlations of these measures, the estimates of the direct effects of age on citations are essentially the same here as in Table 4. The squared deviations of the sentiment measures from the prevailing norms are only weakly statistically significant, with bigger departures from the norm of positivity increasing citations, but with larger decreases in citations among articles whose style departs from norms more along the dimensions of certainty and contemporaneity.

Does writing in a style that departs further—in either direction—from that of other scholars lead to more or less eventual scholarly impact? Taking the pooled estimates (bottom panel of Table 5), we calculate the effect of simultaneous two standard-deviation changes in each measure of sentiment on an article's citations. These departures generate a net reduction of 8 citations, a two-percent drop, but equivalent to only 0.01 standard deviations of citations. Departures in either direction from all three style norms reduce citations, but not by very much. Nonetheless, it is disconcerting that other economists pay less attention to research that is written up in ways that are less familiar.[15]

While doubling the number of authors on an article does not double its citations, it does increase them (Hamermesh, 2018). Since we showed before that co-authorship hardly changes the impact of age on writing style, failure to include the number of authors in these estimates will not bias the estimated impacts here. Adding the number of authors to the models presented in Tables 4 and 5 thus barely alters the results, with some estimates rising slightly in absolute value, some falling, and with those that are statistically significant in Tables 4 and 5 remaining so.[16]

Another potential problem is that more senior authors are more likely to have published more articles included in the sample. If so, and if having published more articles makes additional articles better cited, either because of reputational effects or simply because those who publish more top-level articles do more important work, the estimated effects shown in Tables 4 and 5 may be biased. Age at publication and

[15]In terms of the epigraph to this article, this result suggests that they do bring forth fruit in old age, but that it is not so succulent as the fruit that they brought forth earlier.

[16]Restricting the samples, first to those with more than five entries, then to those with ten or more entries, also does not qualitatively alter the results. Even for the 1970s cohort, for which the second restriction cuts the sample to only 2,167 observations, the parameter estimates retain their signs, and the statistically significant negative estimates for CERT in Table 4, and CONP in Table 5, nearly double in value.

number of articles are correlated but not very highly—r = 0.11 in the entire sample, r = 0.15 in the 1970s cohort.

To examine this possibility, we re-estimated the models in the bottom panels of Tables 4 and 5, adding for each observation the number of prior articles the author wrote that are included in the sample. The estimated effects of z^{*}_{ij}, and z^{*2}_{iaj} on citations do not change very much, with all of them increasing slightly in absolute value from those shown in the Tables. The estimated impact of age on citations is unsurprisingly reduced, by nearly half in some of the estimates. Most interesting, and relevant for the next Section of this study, among authors of the same age, those who had previously published more in these journals receive more citations to their current publication than otherwise identical other authors. We cannot determine whether this treatment reflects higher-quality work or reputational ("Matthew") effects (Merton, 1968). Suffice it to note that the negative impact of age on citations is reduced for authors who are the more successful among the highly successful scholars in this sample.

We can decompose the total effect of age on citations using the estimates in the middle panel of Table 2 and the pooled estimate at the bottom of Table 4 as:

$$(3) \quad dCITS/d_{AGE} = \partial CITS/\partial AGE \,\Big|_{z^{*}ij} + [\partial CITS/\partial z^{*}_{ij} \,\Big|_{AGE} \cdot \partial z^{*}_{ij}/\partial AGE],$$

the sum of the direct effect, the first term, and the indirect effect, the term in brackets. We calculate the effect on citations of a two standard-deviation increase in age in the whole sample (the 1970s cohort). As a fraction of mean citations, the effects are reductions of 3.8 (7.5) percent, which equal 0.01 (0.02) standard deviations in citations (with only two percent of the impact working through the indirect effect in (3)). Scholarly recognition decreases with author's age, but only a minute part of the decrease is due to changes in writing style with age.

We can only speculate about why there are fewer citations to articles published at the same time and in the same sub-field by older scholars and why they decrease as scholars age. One possibility is that already-established authors are favored by editors, who publish their papers even if the work is not quite so important as that of more junior authors or their own earlier work (although some evidence points against this kind of favoritism in several related dimensions, Blank, 1991).[17] No doubt other, perhaps even testable explanations are consistent with this surprising finding.

We can replace $z^*_{ij/}$ by z^{*2}_{iaj} in (3) to calculate the effects of any departures from norms, using the estimates in the middle panel of Table 3 and at the bottom of Table 5. The decomposition differs little from that above, with a slightly larger total effect in the entire sample and a slightly smaller effect for the 1970s cohort. Again, the indirect effects constitute no more than two percent of the total impact.

V. Exits

In this section we examine the patterns of decline in publication with age (shown in Figures 1a and 1b), considering how rates of slowdown relate to prior productivity, to the scholarly impact of prior work, and to the style in which that work is written. We also examine how prior productivity and style relate to exits from academe in the form of retirement, and, as a placebo test, to death. Unlike in the previous sections, where the units of observation were articles, here they are the scholars whose works were included in the previous analysis.

We collected information on each of the authors in the sample, taking from their CVs information on whether by 2018 (and when) they retired, died, or switched out of a career typical among highly successful scholars. In describing the time paths of publications, we use the 1970s, 1980s and 1990s cohorts, 945 of whose 960 members had available CVs, and of whom 78 percent were in academe in 2018. In describing retiring/dying we concentrate on the 1970s cohort. Of the 359 authors in that group, we found CVs of 346. Of them, in 2018 56 percent remained in academe, 27 percent had retired, 5 percent had died. In both sets of analyses, we exclude the small percentages of the samples who had left academia.

[17]Testing this idea by including an interaction with an indicator of whether the author had recently published in a particular journal shows, if anything, that an article is better-cited (although not significantly) if s/he has published recently in the journal. This is not consistent with editors publishing relatively inferior papers by authors whom they had published before.

A. Slowing Down

We estimate a series of autoregressions describing output in each of several post-Ph.D. decades by prior publications and their characteristics:

(4) $A_d = \Sigma b_{1,d-t} A_{d-t} + \Sigma\, b_{2,d-t} CITS^{*}_{d-t} + \Sigma\Sigma b_{3,d-t} z^{*}_{iaj}$, $j=1,2,3$,

where A is the number of articles published in decade d (d=10-19, 20-29, etc.), and t is the length of the lag (in decades). $CITS^{*}$ is the average adjusted citations to the person's three most recent articles before decade d, z^{*}_{iaj} are the average sentiment scores in the person's three most recent articles before decade d, and the b are parameters to be estimated. [18] We also re-estimate (4) replacing the z^{*} by the z^{2*}.

Table 6 shows OLS estimates of (4) for first-order autoregressions only, since higher-order terms in full versions add little to the explanatory powers of the models.[19] (We present estimates of the fully-specified models in Appendix Table C1.) The odd-numbered columns include the vector of the z^{*}, the even-numbered columns include the vector of the z^{2*}. The samples are restricted to those members of the Ph.D.-age cohorts who remained in academe as of 2018—who had not retired, died, or switched occupations by then.[20] For the vectors of average sentiment scores, we present the p-value of the F-statistic jointly testing the constraint that all three sentiment scores, computed as the averages of the most recent three articles' scores, have no impact on the outcomes. In each equation we also include the *JEL* group of an author's most recent publication and the year each author's Ph.D. was received.[21]

[18]Because for some articles we had Web of Science citations, while for others we had Google Scholar citations, to facilitate the estimates we created $CITS^{*}$ equaling Web of Science citations (if available), or 0.294·GS citations, where the multiplier reflects average citations from the two sources in our data. Because we averaged CITS* over three articles, and because we could not obtain citations for a small part of the sample, the sub-sample used in estimating (4) is reduced.

[19]The correct Poisson estimates of these equations imply the same conclusions as the OLS estimates in Table 6.

[20]We impose this restriction because those who are dead, most of those who are retired, and even many of those who have left academia, face different publishing incentives than those who remain academics. In any case, of the original sample 79 percent are included in the estimates in Columns (1) and (2), 71 percent in Columns (3) and (4), and 56 percent in Columns (5) and (6).

[21]Because the fractions of women in the samples of authors with 20+ years of experience are so tiny, we do not include a gender indicator in the estimates in Table 6. Re-estimating all the equations excluding the few women changes no parameter estimate by more than one in the second significant digit.

Columns (1) and (2) estimate the determinants of output in the second decade of these scholars' careers. The autoregressive parameter on output in the first post-Ph.D. decade is only 0.31, reflecting the well-known tapering off of top-level scholarly publication with experience. Citations matter too: Given the number of publications in the first decade of output, more is published in the second decade if the author's most recent publications are better-cited. With CITS* averaging only 93 for articles published by authors in this sample in their first decade, the impact of better-cited work on subsequent output, although statistically significant, is small, with additional publications in the second decade equaling 0.06 standard errors of $A_{10\text{-}19}$ in response to a two standard-error increase in average citations to articles from the first decade. At the extreme of $CITS^{*}_{d-1}$, 1,747, the implied increase in second-decade output is substantial, 0.80 standard errors of $A_{10\text{-}19}$ compared to the average.

The F-statistics testing the joint significance of the estimates of the impacts of the z^{2*} in Column (2) show that these average sentiment scores have no significant effect on publication rates in scholars' second decade. The F-statistic on the estimates in Column (1), however, shows that the direction of style—its positivity, certainty, and temporal orientation—does have significant impacts on subsequent output, with increases in positivity and decreases in certainty significantly related to subsequent publication. Two standard-error increases in positivity and contemporaneity coupled with a similar decrease in certainty are associated with 0.39 standard error more articles in the second decade. As with citations early in the career, style matters; and the nature of the style variables suggests that these impacts may be interpreted as causal.[22]

The estimates in Columns (3)-(6) of Table 6 show that neither citations nor style in the second or third decades of publishing are related to the quantity of subsequent top-level publications (in the third or

[22]Neither the point estimates nor the F-statistics change much if we restrict the sample to authors with three or more entries in their first nine years. With the smaller sample size, the standard errors become almost exactly proportionately larger. The estimated autoregressive parameter increases because of the sample restriction, but still remains smaller than those shown in Columns (3)-(6).

fourth and fifth decades).[23] All that matters is the quantity of output in the previous decade; and it matters more than it does for output in the second: The autocorrelation coefficients increase as careers progress. Whether these changes result from authors' habits becoming more important as they age, or whether reputational effects and editorial inertia are generating them, cannot be inferred from the data—the results cannot be interpreted as solely the results of authors' behavior.

Given the tremendous growth of co-authorship in economics, perhaps co-authorship helps the most successful senior economists maintain publishing at the highest levels (although at lower levels than earlier in their careers). To explore this possibility, we add to each model in Table 6 the number of authors on the person's final paper in decade d-1. These additions produce only minute changes in the estimated autoregressive parameters. Moreover, the impact of recent additional coauthors on subsequent numbers of publications was negative, although never anywhere nearly statistically significant. The few very senior economists who maintain a top-level publication record do not do so by attaching themselves to coauthors.

We specified equation (4) so that the autoregressive parameters do not vary with A_{d-1}. To test this assumption, in additional estimates we replaced the A_{d-1} by vectors of several indicators of the number of publications in the previous decade (e.g., for A_{0-9}, three to five, or more than five publications, with zero to two publications as the base group). The estimates describing publication in the second decade of a career, which were shown in Columns (1) and (2) of Table 6, do change: Authors with three to five publications in their first decade publish less in the second decade than authors who had zero to two publications in their first decade, while those with more than five early publications produce still more in their second decade than in their first.[24] Scholars who are only moderately successful (by the high

[23]The vectors of indicators of the *JEL* group of the most recent publication in decade d-1 are never statistically significant.

[24]This re-specification does not alter the conclusions that citations to articles published in the first decade of a career are positively related to those in the second decade, nor that the style of articles published in the first decade also matters. These results do not depend on the inclusion of a few people whose entire *oeuvre* in the data was produced in their first post-Ph.D. decade: Excluding them from these re-specifications does not change the inferences about the relation between early and subsequent publications.

standards for inclusion in this sample) early on fade, while stars early in their careers become superstars. Regardless, as the re-specifications of the models in Columns (3)-(6) show, even superstars fade; and the estimated parameters on the indicators included in these re-specifications for publications in the third decade or beyond show that the autoregressive parameters in Columns (3)-(6) are not functions of A_{d-1}—that the linear models presented in the Table describe the data well.

B. Stopping

How do a relative lack of recent publishing success, the attention paid to recent research, and its style induce distinguished senior scholars to retire from academe? We restrict the analysis to members of the 1970s cohort, since most authors from later cohorts were too young to have been contemplating retirement by 2018. Because mandatory retirement might affect academics' choices about retiring (not in the U.S., where the majority of members of this cohort worked, but perhaps elsewhere), we collected information on whether and when a scholar would at least nominally have been subject to such a rule, creating an indicator variable for it. We estimate probits on whether the person had retired by 2018, including $A_{20\text{-}29}$, adjusted citations to the three most recent publications before Ph.D. age of 30, vectors of style measures (the z^{*}_{iaj} or z^{2*}_{iaj}), and the *JEL* group of the most recent article before that age, including in the sample all those who remained in academe or were retired in 2018 (leaving codes 1 and 2).

The estimates are presented in Columns (1) and (2) of Table 7. (As with the descriptions of slowing down, shown in Appendix Table C1, Appendix Table C2 shows that additional lags in A had small effects on the probability of retirement and were not statistically significant.) Regardless of which vector of style measures, the z^* the z^{2*}, is included, having published more top-level articles in the third decade of a career leads to a significantly lower likelihood of subsequent retirement from academe. The impact is also not small: Comparing the 36 percent of authors who published no top-level papers in the third decade to those who published four papers then (the 91st percentile in this cohort), the former are 17 percentage points more likely to have retired by 2018 (on a mean retirement probability of 0.32).[25] Moreover, replacing the

[25]In some specifications we added a vector of indicators of the year when a member of this cohort received her/his Ph.D. The estimates were nearly identical to those shown in Table 7, Columns (1) and (2).

continuous measure of recent publications with a set of indicators of the number of publications demonstrates that the negative effect of additional publication on the probability of retirement is essentially linear in A_{20-29}.

If a scholar's recent article is more heavily cited, s/he is less likely to choose to retire, although the estimate is not statistically significant (and the implied impact of additional citations is small). Also, the style of recent publications has no effect on the choice to retire. The possibility of being subject to mandatory retirement nearly doubles the likelihood of being retired in this sample, although the prediction is not perfect (presumably because the laws and other mandates can be circumvented). The main conclusion from these estimates is that what matters for retirement is the quantity of top-level output: Since retirement is the scholar's own choice (although we recognize that demand-side effects, in the form of offers of "golden handshakes," might also affect retirement decisions), this result suggests that the inability to publish at the very highest level of scholarship makes scholars more likely to end their careers.[26]

We can enlarge the sample to 392 observations by adding scholars from the 1960s cohort (Ph.D. between 1959 and 1968) to the estimates and re-estimating the model (adding an indicator of the decade when the Ph.D. was received). The impact of A_{20-29}, becomes slightly larger and more significant statistically (-0.057, s.e. = 0.016) with this expansion, but little else changes. The results in Columns (1) and (2) do not result from restricting the sample to the 1970s cohort, the only one almost all of whose publications are included in our data set and who might be approaching a usual retirement age.

The estimates might simply be an artifact reflecting that the scholars who did retire are different from others in ways that we cannot measure but that are correlated with recent success in publishing. To examine this possibility, we estimate the same equation but with the outcome being whether the scholar had died by 2018, using the sample of those who were either dead or still alive in academe and not retired

[26]Their careers are finished after retirement, at least as measured by top-level publishing. Accounting for possible three-year lags from production of an article to its publication in these journals, those who had retired by 2018 produced only two articles in total more than three years after their retirement date (compared to a total of 775 articles that the retirees had produced before then).

(leaving codes 1 and 4). In Columns (3) and (4) of Table 7 we show the estimates of what might be viewed as a placebo test of the model describing retirement. Although the coefficients on both variables describing recent publications have the same negative signs as in the first two columns, neither is anywhere nearly significant statistically, and both are much smaller in absolute value. Moreover, given evidence that an author's death reduces citations to previously published papers (Aizenman and Kletzer, 2011, and by inference, Azoulay *et al*., 2010) the estimated impact of prior citations in this placebo is probably biased negatively. The same may be true for $A_{20\text{-}29}$ if those who died were relatively unhealthy and perhaps hence less productive during their third decade.[27] These differences suggest that the results on retirement do not arise from correlations of unobservables with both the incidence of retirement and recent publication success.[28]

VI. Conclusions and Speculations

Using analyses of textual styles of 50 years of economics research papers in five major journals, coupled with information on the articles' subsequent citations and their authors' demographic characteristics, particularly their age, we have shown that departures of writing style from contemporaneous norms within a subfield generate less scholarly attention to an article. These deviations increase with authors' ages, contributing a small part of the rather startling decline in attention to articles that are produced by older authors, that we also document. The rate of slowdown in publishing with age is a linear function of an author's prior productivity, but the rate of slowdown in mid-career is greater if an author's prior work has been less well-cited. Having produced less top-flight research late in a career leads people to choose to retire earlier.

We stress that all of our conclusions are based on a sample of the very top scholars in economics, and that we cannot infer from this selected sample whether similar changes with age occur in the careers of

[27]We also estimate a multinomial logit for members of both samples used in this sub-section, with the outcomes remaining in academe, retiring, or dying. The estimates are almost identical to those listed in Table 7.

[28]We can also interpret these results as a direct test of "publish or perish." In this sample the estimates suggest that these are not alternatives.

less successful Ph.D. economists in the same cohorts. With that caveat we have documented one source of the diminution of top-flight scholarly activity with age—the decreasingly warm reception paid by other scholars, due in small part to changes in the style of writing as an author ages. But other than randomness, the main apparent cause of the decline is habit: Those scholars who have been most productive remain so, albeit at a diminished rate of productivity. "Pooping out" is mostly endogenous, whether because of technological obsolescence, loss of interest (one's own or that of editors), boredom, reduced financial incentives, flight to administrative roles or other alternative paths, or some other factor or combination of factors.

The various findings suggest a variety of additional questions, some of which might be answerable with additional data. Here we can only speculate about them in the context of our results. For example, academic economics for those near the top of their field, like the people in our sample, is a very easy existence: Minimal teaching burdens, no publication requirements, and salaries that may not increase with academic pay generally but that are far above average pay in an economy. Why retire? Is it a desire for uninterrupted leisure, including the complementarity of non-work time and substantial retirement pay; embarrassment at being unable to keep up with the publication success and interest/enthusiasm of younger colleagues; financial bonuses that induce retirement? Using academia as an example for high-paying occupations should be a way to learn more about why people generally retire rather than stay on or switch to part-time work (which academic jobs can *de facto* if not *de jure* become).

Future research exploring these various motivations would be particularly useful from a policy angle. Understanding the causes of declining output with age among top academic researchers might lead to the construction of appropriate financial incentives, technological assistance, or some other such malleable factor, that could keep top-level output continuing. Findings on this specific issue might even apply not just to academia, but to other fields where top-level employees' productivity tends to decline with age. For sectors and industries that are "aging out," or having difficulty attracting younger workers, this could be very important indeed.

REFERENCES

Joshua Aizenman and Kenneth Kletzer, "The Life Cycle of Scholars and Papers in Economics: The 'Citation Death Tax'," *Applied Economics*, 43 (Nov. 2011): 4135-48.

Orley Ashenfelter and David Card," Did the Elimination of Mandatory Retirement Affect Faculty Retirement," *American Economic Review*, 92 (Sept. 2002): 957-80.

Pierre Azoulay, Joshua Graff Zivin, and Jialan Wang, "Superstar Extinction," *Quarterly Journal of Economics*, 125 (May 2010): 549-89.

Rebecca Blank, "The Effects of Double-Blind versus Single-Blind Reviewing: Experimental Evidence from the *American Economic Review*," *American Economic Review*, 81 (Dec. 1991): 1041-67.

George Borjas and Kirk Doran, "Prizes and Productivity: How Winning the Fields Medal Affects Scientific Output Prizes and Productivity," *Journal of Human Resources*, 50 (Summer 2015): 728-58.

Daniele Checchi, Alberto Ciolfi, Gianni De Fraja, Irene Mazzotta, and Stefano Verzillo, "Have You Read This? An Empirical Comparison of the British REF Peer Review and the Italian VQR Bibliometric Algorithm," *Economica*, 88 (Oct. 2021): 1107-29.

Tom Coupé, Valerie Smeets, and Frédéric Warzynski, "Incentives, Sorting, and Productivity long the Career: Evidence from a Sample of Top Economists," *Journal of Law, Economics and Organization*, 22 (April 2006): 137-67.

Arthur Diamond, "What Is a Citation Worth," *Journal of Human Resources*, 21 (Spring 1986): 200-15.

Matthew Gentzkow, Bryan Kelly, and Matt Taddy, "Text as Data," *Journal of Economic Literature*, 57 (Sept. 2019): 535-74.

Amanda Goodall, *Socrates in the Boardroom: Why Research Universities Should Be Led by Top Scholars.* Princeton, NJ: Princeton University Press, 2010.

Daniel Hamermesh, "Citations in Economics: Measurement, Uses, and Impacts," *Journal of Economic Literature*, 56 (March 2018): 115-56.

----------------------, George Johnson, and Burton Weisbrod, "Scholarship, Citations, and Salaries: Economic Rewards in Economics," *Southern Economic Journal*, 49 (Oct. 1982): 472-81.

--------------------- and Sharon Oster, "Aging and Productivity among Economists," *Review of Economics and Statistics*, 80 (Feb. 1998): 154-6.

--------------------- and Gerard Pfann, "Reputation and Earnings: The Roles of Quantity and Quality in Academe," *Economic Inquiry*, 50 (Jan. 2012): 1-16.

James Heckman and Sidharth Moktan, "Publishing and Promotion in Economics: The Tyranny of the `Top 5.'" *Journal of Economic Literature* 58 (June 2020): 419-70.

Alphonse G. Holtmann and Alan Bayer, "Determinants of Professional Income among Recent Recipients of Natural Science Doctorates," *Journal of Business*, 43 (Oct. 1970): 410-18.

Cornelius Houk, “Statistical Analysis of Genesis Sources,” *Journal for the Study of the Old Testament*, 27 (Sept. 2002): 75-99.

Anthony Kenny, *A Stylometric Study of the New Testament*. New York: Oxford University Press, 1986.

Lea-Rachel Kosnik, “Who Are the More Dismal Economists? Gender and Language in Academic Economics Research,” American Economic Association, *Papers and Proceedings*, 112 (May 2022): 592-6.

Harvey Lehman, *Age and Achievement*. Princeton, NJ: Princeton University Press, 1953.

Sharon Levin and Paula Stephan, “Research Productivity over the Life Cycle: Evidence for Academic Scientists,” *American Economic Review,* 81 (March 1991): 114-32.

Colin Martindale, *The Clockwork Muse: The Predictability of Artistic Change*. New York: Basic Books, 1990.

Robert Merton, “The Matthew Effect in Science,” *Science,* 159 (Jan. 5, 1968).

Frederick Mosteller and David Wallace, “Inference in an Authorship Problem,” *Journal of the American Statistical Association*, 58 (June 1963): 275-309.

Andrew Oswald, “An Examination of the Reliability of Prestigious Scholarly Journals: Evidence and Implications for Decision-Makers,” *Economica,* 74 (Feb. 2007): 21-31.

James Pennebaker, Ryan Boyd, Kayla Jordan, and Kate Blackburn, *The Development and Psychometric Properties of LIWC2015*. Austin, TX: University of Texas at Austin, 2015.

Sascha Schweitzer and Jan Brendel, “A Burden of Knowledge Creation in Academic Research: Evidence from Publication Data,” *Industry and Innovation*, 28 (2021): 283-306.

Bruce Weinberg and David Galenson, “Creative Careers: The Life Cycles of Nobel Laureates in Economics,” *De Economist,* 167 (2019): 221-39.

Table 1. Descriptive Statistics of the Sample of Articles, Authors, and Sentiments.

Articles (N=9,280)

JEL Group:	**%**	**Decade**	**%**	**Pages—Mean (s.d.)**
Theory and methodology	22.1	1969-78	16.7	11.92 (6.80)
Microeconomics, industrial organization	25.9	1979-88	22.7	14.41 (6.62)
Macroeconomics, international, financial	25.6	1989-98	19.4	19.80 (7.84)
Public, health/education, labor	15.3	1999-2008	21.0	25.79 (8.96)
Other	11.1	2009-18	20.2	34.30 (11.02)
	100.0		100.0	

Authors

		Entire sample (N=1,389)	**1969-1978 cohort (N=359)**
		%	**%**
N articles:	5	24.2	22.6
	6-10	49.5	45.4
	11-20	21.5	25.9
	21+	4.8	6.1
		100.0	100.0

Sentiment –Mean (s.d.) Articles (N=9,280)

	Raw	**Deviation**	**Deviation^2**
POSN	-0.229 (0.189)	0.003 (0.185)	0.034 (0.049)
CERT	-0.358 (0.171)	-0.008 (0.170)	0.029 (0.051)
CONP	0.723 (0.134)	-0.0003 (0.130)	0.017 (0.048)

Table 2. Relationship of Relative Style to Ph.D. Age[a]

	Entire sample, 12,814 articles, 1,389 authors				1970s Cohort, 3,562 articles, 359 authors		
Dep. Var./100 :	*POSN*	*CERT*	*CONP*		*POSN*	*CERT*	*CONP*
				Simple regression[b]			
Years past Ph.D.	0.0697	-0.0408	-0.0609		0.0465	-0.0184	-0.0430
	(0.0245)	(0.0222)	(0.0243)		(0.0439)	(0.0420)	(0.0390)
R^2	0.002	0.001	0.002		0.001	0.001	0.004
				Adds decade, journal, gender, *JEL* group[c]			
Years past Ph.D.	0.0770	-0.0531	-0.0789		0.0597	-0.0340	-0.0183
	(0.0282)	(0.0258)	(0.0272)		(0.0455)	(0.0429)	(0.0352)
R^2	0.013	0.015	0.048		0.025	0.013	0.063
				Author fixed effects			
Years past Ph.D.	0.1408	-0.0492	-0.0661		0.0291	-0.0234	-0.0308
	(0.0580)	(0.0524)	(0.0353)		(0.0380)	(0.0351)	(0.0228)
R^2	0.286	0.307	0.459		0.299	0.294	0.447
				Mean (s.d.)			
Years past Ph.D.		12.98				12.91	
		(9.48)				(10.28)	

[a]Standard errors in parentheses, clustered on authors.
[b]Includes *AER*-equivalent page count
[c]Gender not included in fixed-effects estimates. Five decades (except decade in cohort estimates), five *JEL* groups, five journals.

Table 3. Relationship of Squared Style Deviation to Ph.D. Age

	Entire sample, 12,814 articles, 1,389 authors				1970s Cohort, 3,562 articles, 359 author		
Dep. Var.:	*POSN*^{2}	*CERT*^{2}	*CONP*^{2}		*POSN*^{2}	*CERT*^{2}	*CONP*^{2}
				Simple regression[a]			
Years past Ph.D./100	0.0029	0.0099	0.0320		0.0452	0.0471	0.0481
	(0.0056)	(0.0055)	(0.0092)		(0.0111)	(0.0121)	(0.0179)
Adj. R^2	0.002	0.011	0.004		0.017	0.026	0.009
				Adds decade, journal, gender, *JEL* group[b]			
Years past Ph.D./100	-0.0058	-0.0023	0.0287		0.0433	0.0475	0.0400
	(0.0062)	(0.0061)	(0.0100)		(0.0116)	(0.0113)	(0.0171)
Adj. R^2	0.014	0.018	0.028		0.022	0.029	0.045
				Author fixed effects			
Years past Ph.D./100	0.0239	0.0670	0.0309		0.0288	0.0482	0.0383
	(0.0162)	(0.0169)	(0.0143)		(0.0109)	(0.0126)	(0.0086)
Adj. R^2	0.193	0.203	0.315		0.212	0.204	0.325

[a]Includes *AER*-equivalent page count
[b]Gender not included in fixed-effects estimates. Five decades (except decade in cohort estimates), five *JEL* groups, five journals.

Table. 4. Relationship of Citations to Relative Style and Ph.D. Age, LAD Estimates[a]

Effect on citations of:	Entire sample, 12,740 articles, 1,389 authors				1970s cohort, 3,531 articles, 359 authors		
	POSN	*CERT*	*CONP*		*POSN*	*CERT*	*CONP*
				Simple LAD[b]			
Sentiment	-35.53	-66.11	-110.46		-18.90	-69.26	-69.24
	(9.03)	(9.79)	(18.07)		(17.34)	(16.95)	(37.64)
Years past Ph.D.	-1.55	-1.59	-1.65		-0.45	-0.63	-0.51
	(0.20)	(0.19)	(0.19)		(0.35)	(0.38)	(0.34)
R^2	0.035	0.036	0.032		0.022	0.024	0.022
				Adds year, journal, gender, *JEL* group[c]			
Sentiment	-28.03	-59.69	-46.29		-4.74	-61.85	-36.38
	(8.46)	(8.97)	(14.03)		(16.58)	(15.80)	(42.44)
Years past Ph.D.	-0.70	-0.69	-0.71		-1.29	-1.51	-1.39
	(0.16)	(0.17)	(0.17)		(0.41)	(0.39)	(0.40)
R^2	0.041	0.041	0.039		0.023	0.025	0.023
				Pooled Sentiments			
Sentiment	-19.60	-57.73	-28.39		-0.17	-66.31	-22.81
	(8.41)	(8.79)	(13.04)		(15.68)	(17.48)	(40.20)
Year past Ph.D.		-0.74				-1.51	
		(0.16)				(0.39)	
R^2		0.041				0.025	

[a]Standard errors in parentheses, cluster on authors.
[b]Also includes *AER*-equivalent page count and an indicator for Google Scholar vs. Web of Science citation data.
[c]Publication year not included in the estimates for the cohort.

Table 5. Relationship of Citations to Squared Style Deviation and Ph.D. Age, LAD Estimates[a]

Effect on citations:	Entire sample, 12,740 articles, 1,389 authors			1970s Cohort, 3,531 articles, 359 authors		
	POSN^2	*CERT*^{2}	*CONP*^{2}	*POSN*^{2}	*CERT*^{2}	*CONP*^{2}
Simple LAD[b]						
Sentiment	-1.95	-111.92	-40.98	20.52	-38.82	-32.83
	(34.69)	(21.39)	(44.52)	(56.38)	(21.07)	(78.03)
Years past Ph.D.	-1.61	-1.63	-1.58	-0.47	-0.43	-0.46
	(0.20)	(0.20)	(0.20)	(0.36)	(0.36)	(0.36)
R^2	0.034	0.034	0.034	0.022	0.022	0.022
Adds year, journal, gender, *JEL* group[c]						
Sentiment	57.79	-67.46	-77.38	22.44	-30.65	-78.01
	(33.45)	(18.26)	(38.84)	(52.93)	(22.02)	(47.35)
Years past Ph.D.	-0.69	-0.71	-0.74	-1.33	-1.31	-1.25
	(0.17)	(0.17)	(0.17)	(0.40)	(0.41)	(0.39)
R^2	0.040	0.040	0.040	0.023	0.023	0.023
Pooled Sentiments^2						
Sentiment	53.30	-56.93	-79.61	26.94	-30.40	-76.24
	(36.29)	(30.78)	(39.12)	(54.96)	(22.04)	(44.30)
Years past Ph.D.		-0.71			-1.24	
		(0.17)			(0.39)	
R^2		0.040			0.023	

[a]Standard errors in parentheses, cluster on authors.
[b]Also includes *AER*-equivalent page count and an indicator for Google Scholar vs. Web of Science citation data.
[c]Publication year not included in the estimates for the 1970s cohort.

Table 6. First-order Autoregressions of Decadal Publications[a]

Ind. Var.:	Decade:	2nd[b]		3rd[c]		4th or 5th[d]	
A_{d-1}		0.317	0.317	0.530	0.526	0.599	0.606
		(0.037)	(0.037)	(0.035)	(0.034)	(0.055)	(0.056)
$(CITES^*_{d-1})/100$		0.053	0.047	0.021	0.018	0.049	0.003
		(0.024)	(0.024)	(0.021)	0.020)	(0.071)	(0.071)
$(z_{ija})_{d-1}$ vector		0.02		0.51		0.09	
$(z_{ija})^2_{d-1}$ vector			0.41		0.32		0.32
(p-value of F(3, N-K))							
R^2		0.160	0.151	0.400	0.402	0.514	0.506
N		715	715	470	470	202	202

[a]Includes indicators for individual year of Ph.D. and JEL group of final article in previous decade.
[b]All authors with Ph.D. year<1999, >1968, who remained in academia.
[c]All authors with Ph.D. year<1989, >1968, who remained in academia.
[d]All authors with Ph.D. year<1979, >1968, who remained in academia.

Table 7. Determinants of the Probability of Exiting Academia After 30+ Years[a]

	Retire[b]		Die[b]	
$A_{20\text{-}29}$	-0.043	-0.045	-0.014	-0.015
	(0.018)	(0.018)	(0.010)	(0.010)
$(CITES^*_{20\text{-}29})/100$	-0.028	-0.028	-0.013	-0.011
	(0.020)	(0.020)	(0.013)	(0.012)
Subject to mandatory retirement	0.261	0.260	0.048	0.043
	(0.090)	(0.089)	(0.073)	(0.071)
$(z_{ija})_{d-1}$ vector	0.96		0.87	
$(z_{ija})^2_{d-1}$ vector		0.43		0.52
(p-value of F(3, N-K))				
Pseudo-R^2	0.079	0.086	0.054	0.068
N	281	281	204	204

[a]Probit derivatives, including indicators for *JEL* group.
[b]Ph.D. year 1969-78, in academia for 30+ years.

Table A1. Examples of Positive and Negative Words in Text

Positive	Negative
optimal	low
satisfy*	bad
good	lack of
efficien*	without
incentive	cannot
consistent	negative
no doubt	work
perfect	poor
unique	no information
improve*	reject*

Table A.2. Examples of Certain and Tentative Words in Text

Certainty	Tentativeness
always	almost
clearly	depending
correct	doubtfully
definitely	generally
every time	might
invariably	sometimes
irrefutably	sort of
truly	suppose
undeniably	unclear
wholly	vaguely

Table A3. Examples of Contemporary and Past Verbs in Text

Contemporary	Past
admit	admitted
arrives	arrived
follows	followed
happens	happened
manage	managed
knows	knew
ranks	ranked
sees	saw
trusts	Trusted
wants	Wanted

Table B1. Journal Style Scores, Adj. for JEL Code and Year, 1969-2018, N = 9,280[a]

	POSN	*CERT*	*CONP*
AER	-------	-------	-------
ETRCA	0.0439	0.0369	0.0565
	(0.0059)	(0.0054)	(0.0040)
JPE	0.0073	0.0094	0.0013
	(0.0060)	(0.0055)	(0.0041)
QJE	0.0043	-0.0136	0.0056
	(0.0061)	(0.0056)	(0.0042)
REStud	0.0448	0.0414	0.0629
	(0.0061)	(0.0056)	(0.0042)
R^2	0.017	0.018	0.047
Range	[-1, 0.45]	[-1, 1]	[-0.529, 1]

[a]Includes indicators for individual years and main *JEL* codes.

Table B2. Correlation Matrices of Journal Style Scores Adjusted for JEL Code and Year, 1969-2018

Sample Period and Size

	1969-2018, N = 12,814			**1969-78 Cohort, N = 3,531**	
	CERT	*CONP*		*CERT*	*CONP*
POSN	0.087	0.128	*POSN*	0.069	0.056
CERT	-------	0.122	*CERT*	-------	0.051
	$CERT^2$	$CONP^2$		$CERT^2$	$CONP^2$
$POSN^2$	0.069	0.057	$POSN^2$	0.064	0.058
$CERT^2$	-------	0.051	$CERT^2$	-------	0.001

Table C1. N^{th}-order Autoregressions of Decadal Publications[a]

Ind. Var:	Decade:	3rd[b]		4th or 5th[c]	
A_{d-1}		0.506	0.508	0.450	0.478
		(0.037)	(0.036)	(0.061)	(0.058)
A_{d-2}		0.053	0.046	0.096	0.114
		(0.033)	(0.032)	(0.053)	(0.051)
A_{d-3}				0.066	0.081
				(0.042)	(0.039)
(Average Cites_{d-1})/100		0.040	0.038	0.055	0.021
		(0.023)	(0.023)	(0.065)	(0.062)
(Average Cites_{d-2})/100		-0.029	-0.030	-0.011	-0.021
		(0.023)	(0.023)	(0.027)	(0.027)
(Average Cites_{d-3})/100				0.039	0.042
				(0.033)	(0.030)
$(z_{ija})_{d-1}$ vector		0.31			
$(z_{ija}^2)_{d-1}$ vector			0.27		
(p-value of F(6, N-K))					
$(z_{ija})_{d-1}$ vector				0.83	
$(z_{ija}^2)_{d-1}$ vector					0.08
(p-value of F(9, N-K))					
R^2		0.417	0.417	0.505	0.535
N		453	453	189	189

[a]Includes indicators for individual year of Ph.D. and JEL group of final article in previous decade.
[b]All authors with Ph.D. year 1969-88 who remained in academia.
[c]All authors with Ph.D. year 1969-78, who remained in academia.

Table C2. Longer Lags in the Determinants of Retirement or Death After 30+Years in Academia[a]

Ind. Var.:	Retire[b]		Die[b]	
A_{20-29}	-0.038	-0.040	-0.018	-0.016
	(0.024)	(0.023)	(0.008)	(0.009)
A_{10-19}	-0.010	-0.009	-0.001	-0.004
	(0.016)	(0.016)	0.004)	(0.005)
A_{0-9}	-0.001	0.002	-0.003	-0.005
	(0.013)	(0.013)	(0.004)	(0.004)
(Average Cites_{20-29})/100	-0.018	-0.033	-0.002	-0.004
	(0.022)	(0.022)	(0.007)	(0.007)
(Average Cites_{10-19})/100	-0.001	0.003	0.0002	0.0004
	(0.010)	(0.009)	(0.003)	(0.004)
(Average Cites_{0-9})/100	-0.010	-0.007	-0.002	0.0002
	(0.012)	(0.012)	(0.003)	(0.003)
Mandatory retirement	0.251	0.228	0.006	0.009
	(0.098)	(0.096)	(0.036)	(0.037)
$(z_{ija})_{d-t}$ vectors	0.53		0.36	
$(z_{ija}^2)_{d-t}$ vectors		0.50		0.39
(p-value of F(9, N-K))				
Pseudo-R^2	0.098	0.100	0.218	0.223
N	262	262	190	190

[a]Probit derivatives, including indicators for *JEL* group.
[b]Ph.D. year 1969-78, in academia for 30+ years.

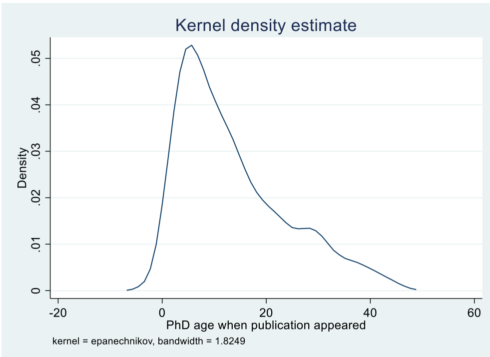

Figure 1a. Kernel Density Estimate of the Distribution of Authors' Ph.D. Ages, Star Authors "Top 5" Journals, 1969-2018.

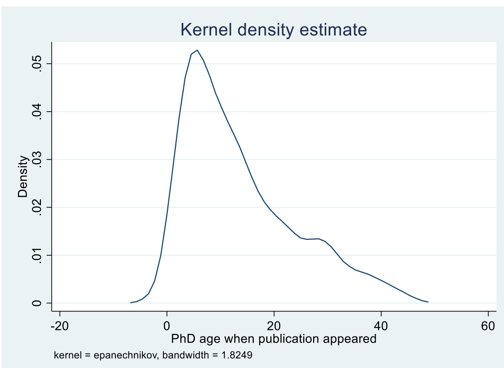

Figure 1b. Kernel Density Estimate of the Distribution of Authors' Ph.D. Ages, Star Authors 1969-78 Cohort, "Top 5" Journals, 1969-2018.

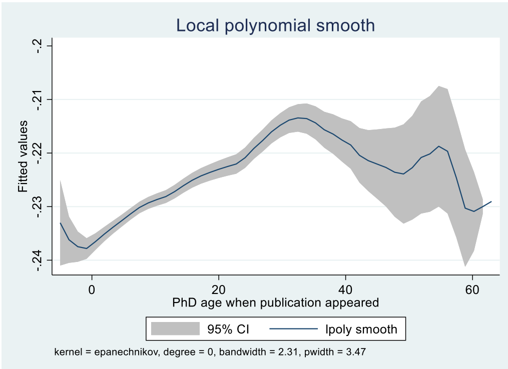

Figure 2a. Local Polynomial Smoothed Relation of Adjusted +/- Sentiment to Ph.D. Age, "Top 5" Journals, 1969-2018 (N = 12,814)

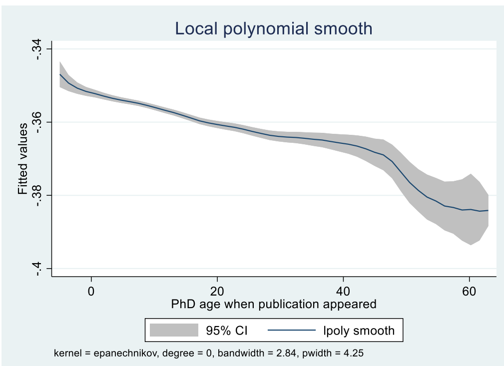

Figure 2b. Local Polynomial Smoothed Relation of Adjusted Certainty to Ph.D. Age, "Top 5" Journals, 1969-2018 (N = 12,814)

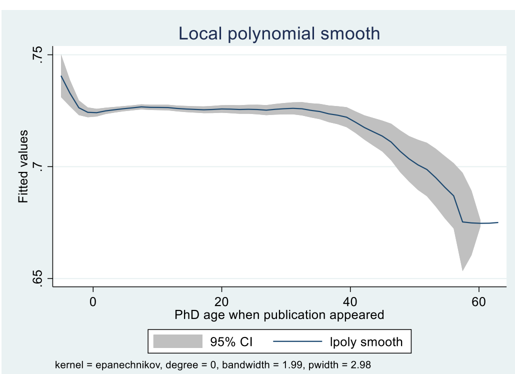

Figure 2c. Local Polynomial Smoothed Relation of Adjusted Present Orientation to Ph.D. Age, "Top 5" Journals, 1969-2018 (N = 12,814)

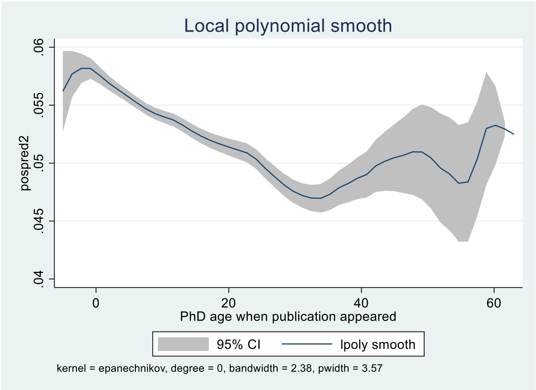

Figure 3a. Local Polynomial Smoothed Relation of Squared Adjusted +/- Sentiment to Ph.D. Age, "Top 5" Journals, 1969-2018 (N = 12,8144)

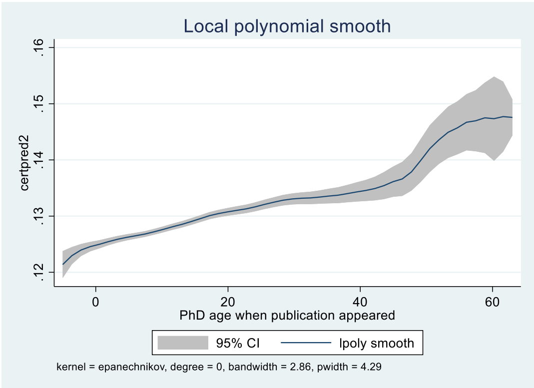

Figure 3b. Local Polynomial Smoothed Relation of Squared Adjusted Certainty Sentiment to Ph.D. Age, "Top 5" Journals, 1969-2018 (N = 12,814)

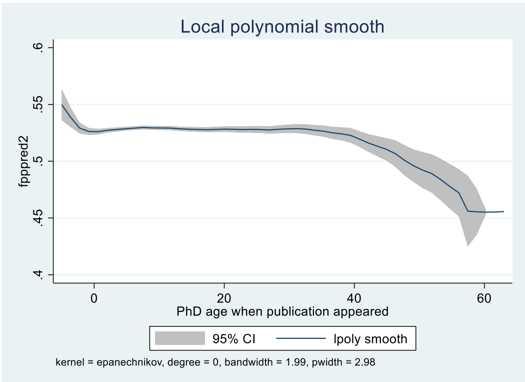

Figure 3c. Local Polynomial Smoothed Relation of Squared Adjusted Present Orientation to Ph.D. Age, "Top 5" Journals, 1969-2018 (N = 12,814)

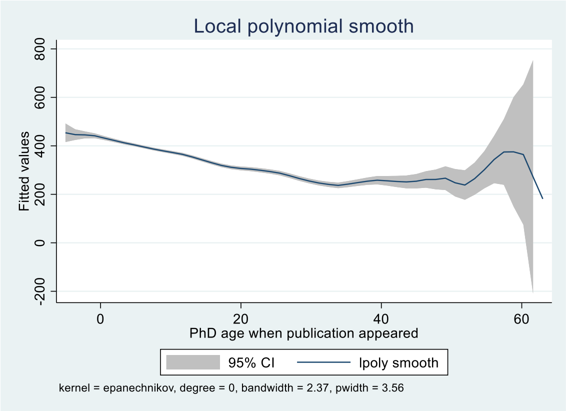

Figure 4a. Local Polynomial Smoothed Relation of Adjusted Citations (for Year of Publication and Citations Measure) to Ph.D. Age, "Top 5" Journals, 1969-2018 (N = 12,740)

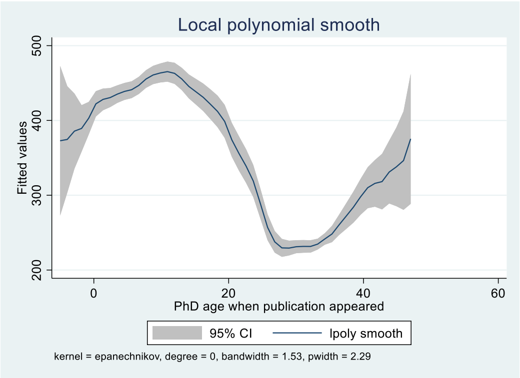

Figure 4b. Local Polynomial Smoothed Relation of Adjusted (for Year of Publication and Citations Measure) Citations to Ph.D. Age, 1969-78 Cohort, "Top 5" Journals, 1969-2018 (N = 3,531)